Collaborative Innovation; A New Lever in Information Technology Development

Bernard Meyerson

IBM Fellow

VP Strategic Alliances and Chief Technologist

IBM Systems & Technology Group, Development

A Question To Ponder:

Where have all the "Gigahertz" gone?

Moore's Law – Economics Made Simple

"Cramming more components onto Integrated Circuits" by Gordon Moore, Electronics, Vol. 38, April 19, 1965.

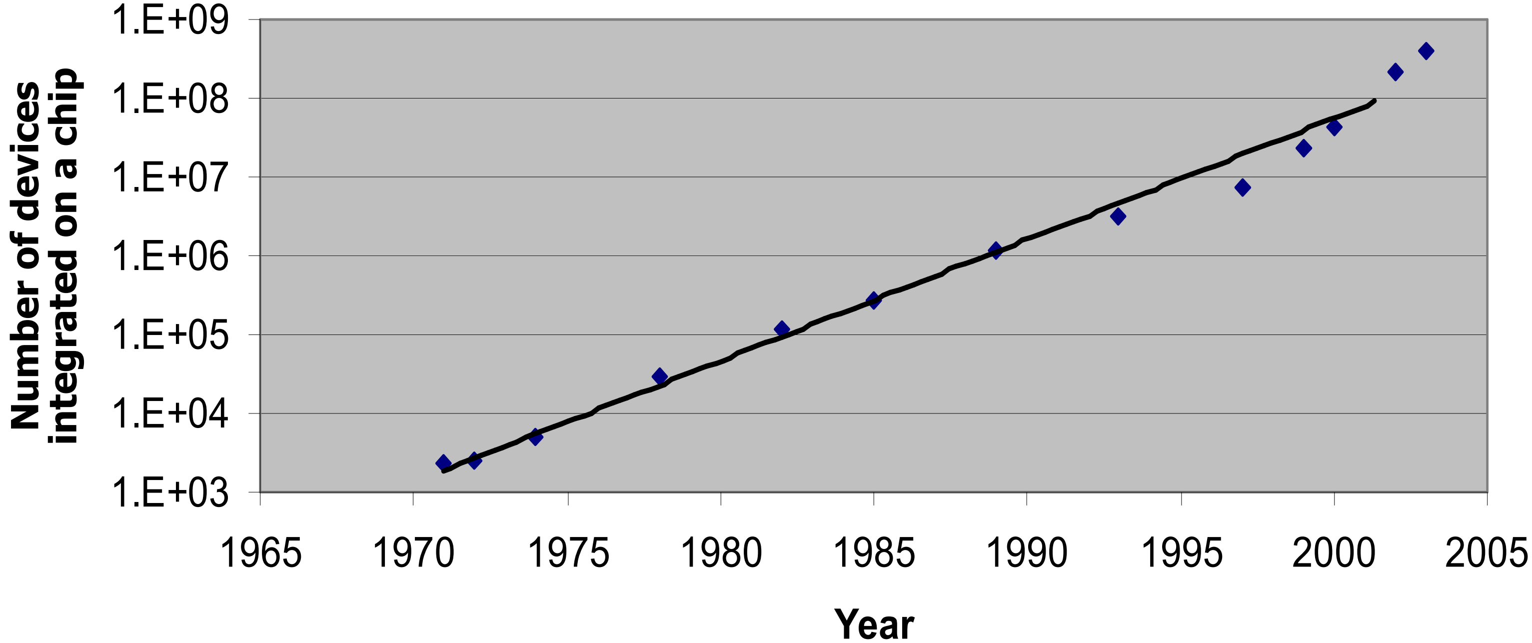

Number of devices integrated on a chip of fixed area doubles every 12-18 months

- Moore's law speaks **ONLY** to the **density** of technology

Classical CMOS Scaling — Is it Really Dead?

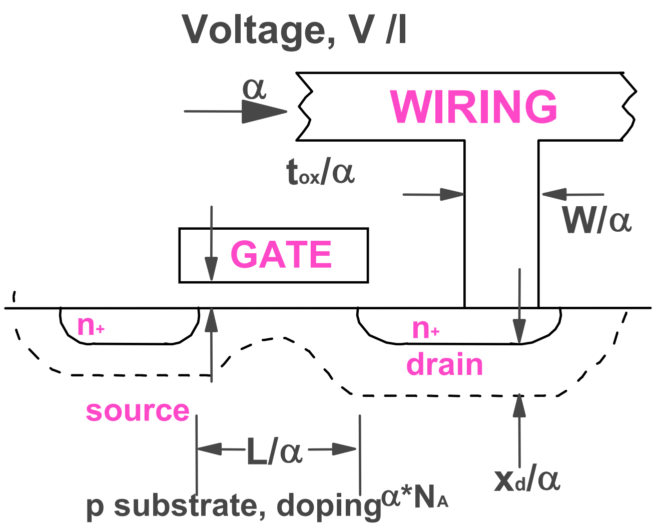

SCALING:

Voltage: V/α
Oxide: t_{ox}/α
Wire width: W/α
Gate width: L/α
Diffusion: x_d/α
Substrate: $\alpha * N_A$

RESULTS:

Higher Density: ~a^2
Higher Speed: ~a
Power/ckt: ~$1/a^2$

Power Density:~Constant

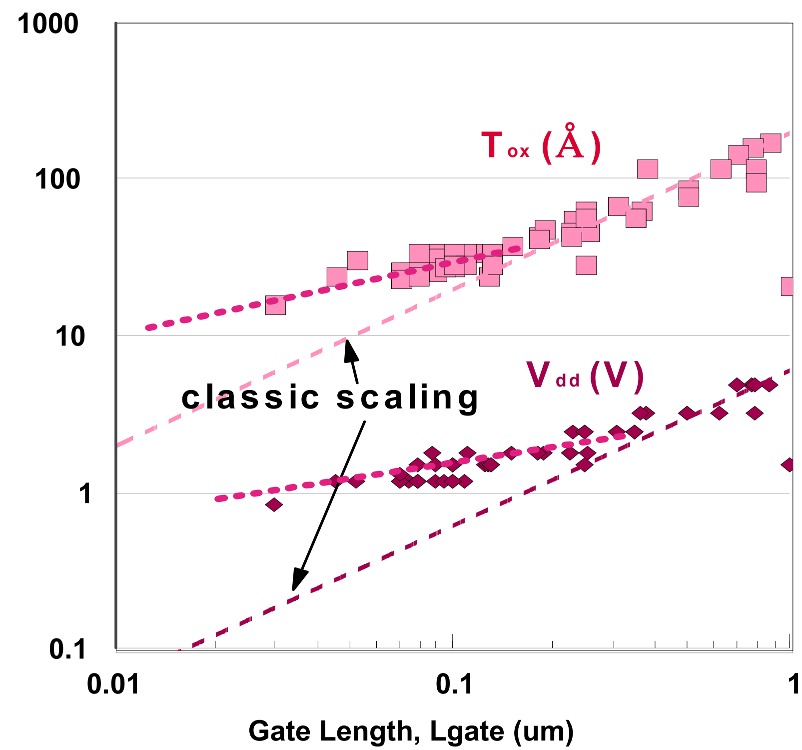

Why deviate from "ideal" scaling
- unacceptable gate leakage/reliability
- additional performance at higher voltage

What is consequence of this deviation?
- **a dramatic rise in power density**

Why Classical Scaling Died: ATOMS DON'T SCALE!

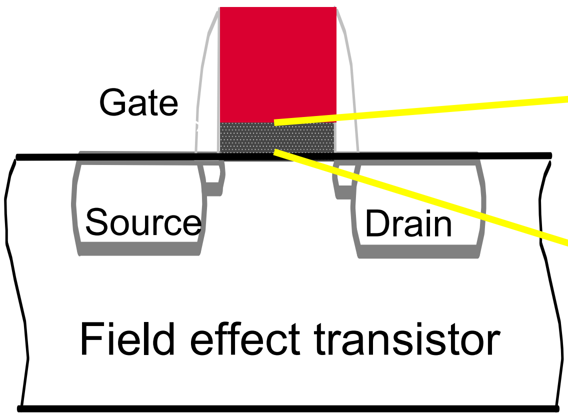

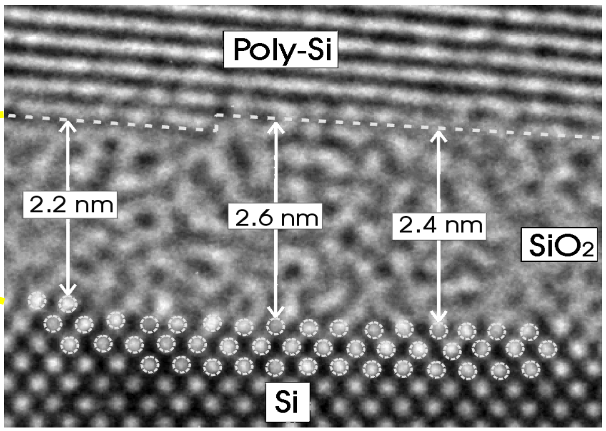

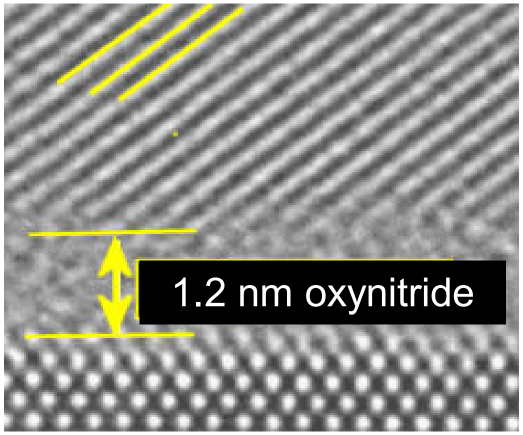

- **Assume only 1 atom high "defects" on each surrounding silicon layer**
  - For a modern "scaled" oxide, ~5-6 atoms thick, 33% variability is induced
- **The <u>bad news</u>**
  - Single atom defects can cause local current leakage 10-100x higher than average
- **The <u>really bad news</u>**
  - Such "non-statistical behaviors" are appearing elsewhere in technology
    - Unless you're a company driving frontier of device & circuit simulation

The Power Cliffs, And a Great Set of Headlights

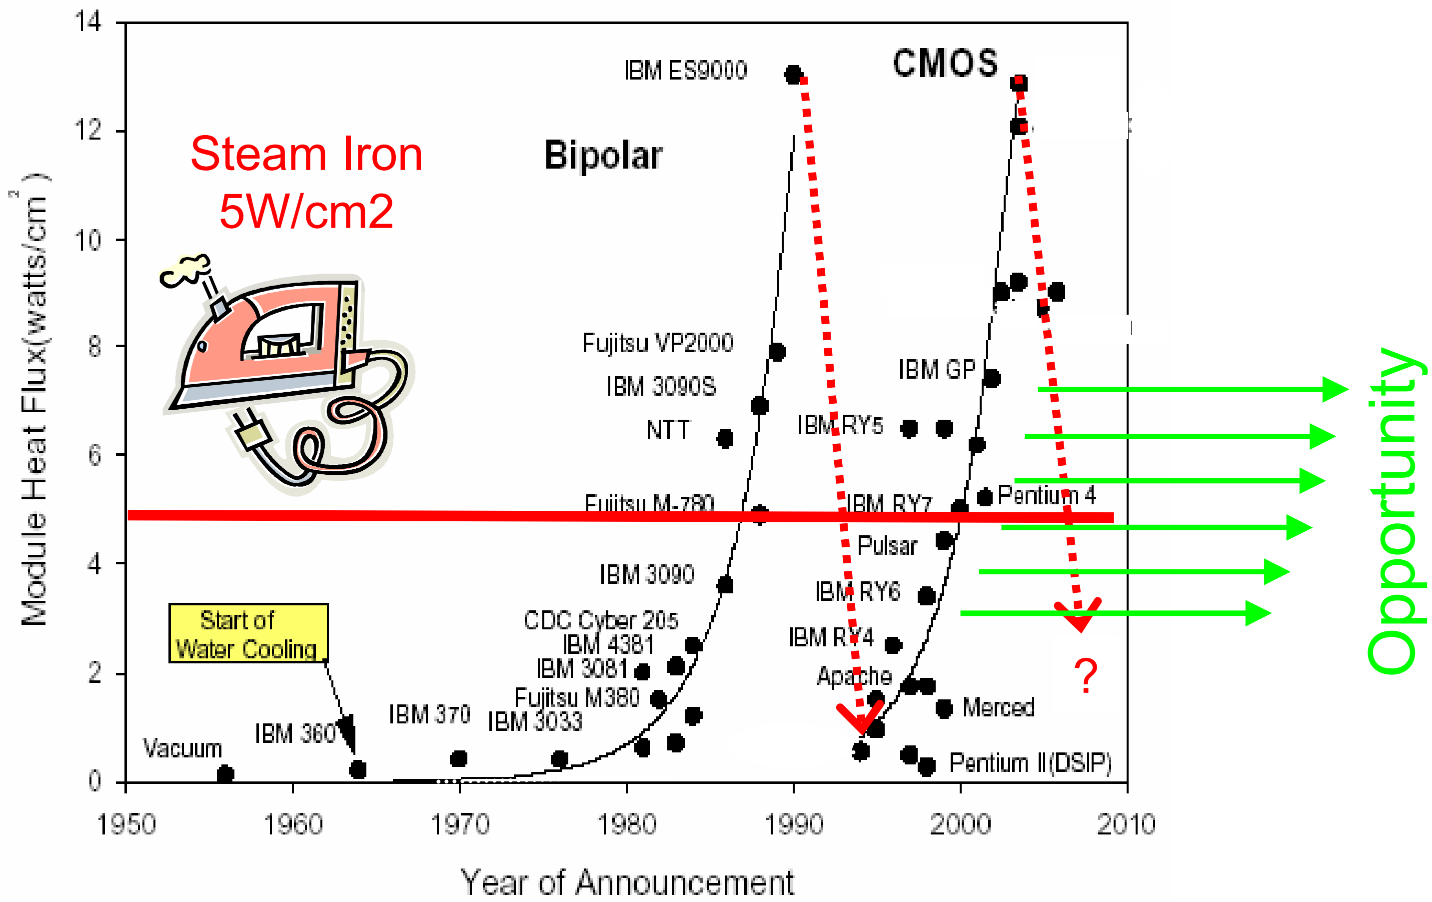

The Discontinuity

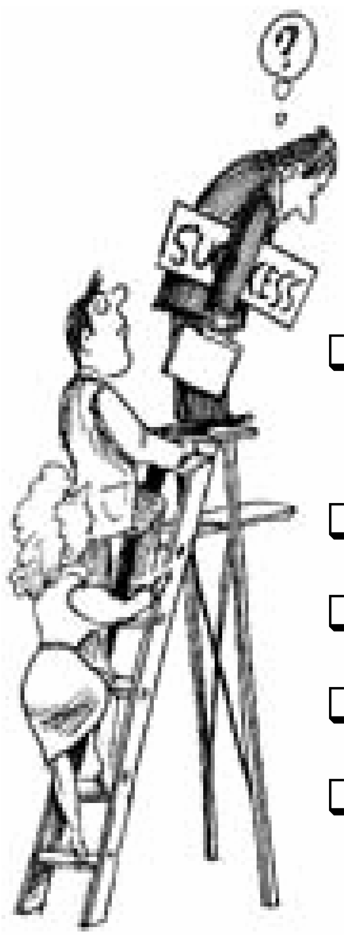

Then (~2003)

- Scaling drove performance
- Scaling drove down cost
- Performance constrained
- Active power dominates
- Focus on **processor** performance

Now

- **Innovation** drives performance
- Scaling drives down cost
- Power constrained
- Standby power dominates
- Focus on **system** performance

Innovation:
Device & Interconnect Enhancements

Novel Materials

Novel Structures

Novel Processes

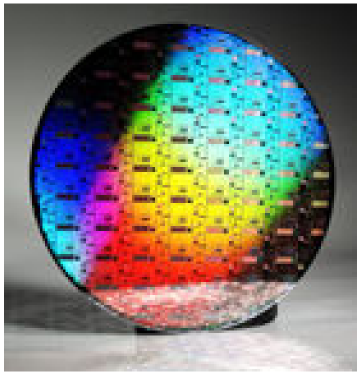

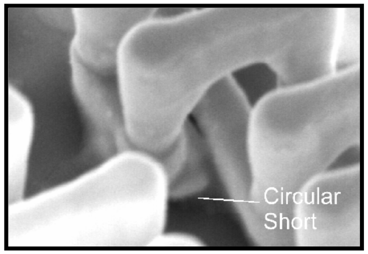

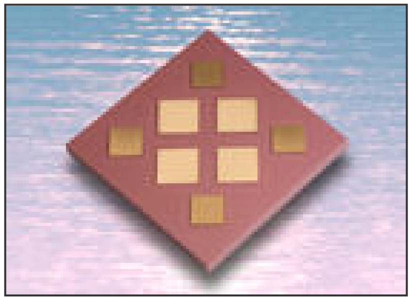

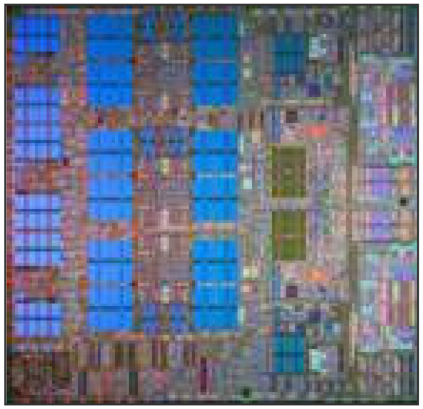

Better Performance Without Scaling

Proven innovation leadership in novel devices & technologies

1994 SiGe

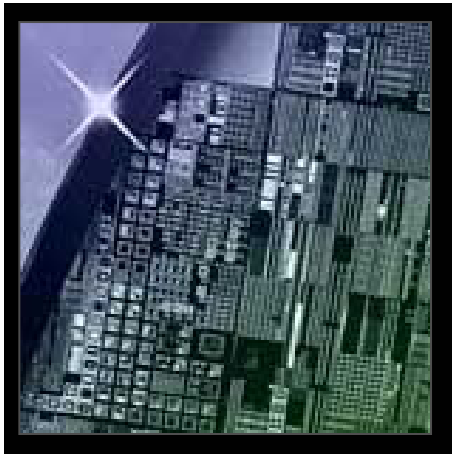

1997 Copper

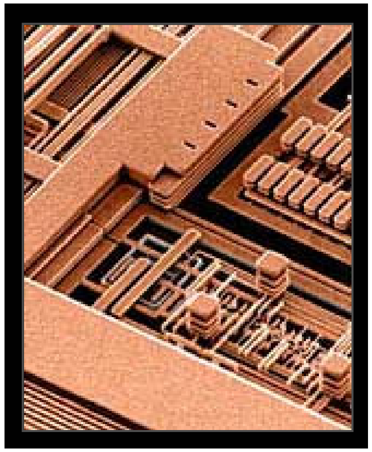

1998 Silicon-on-Insulator

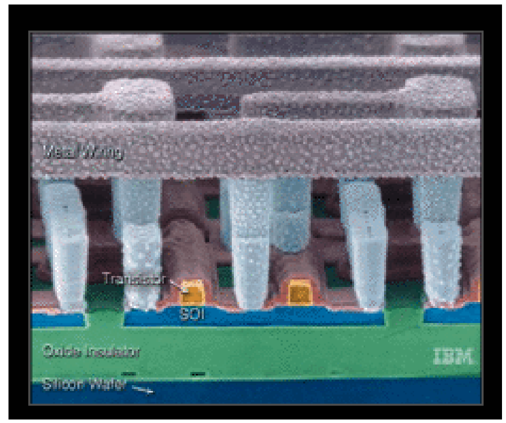

2000 Low-k Dielectric

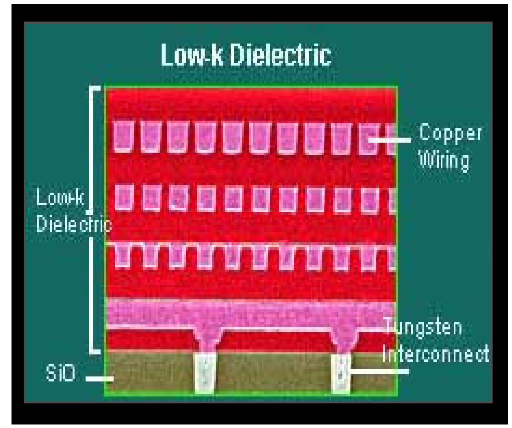

2001 Strained Silicon

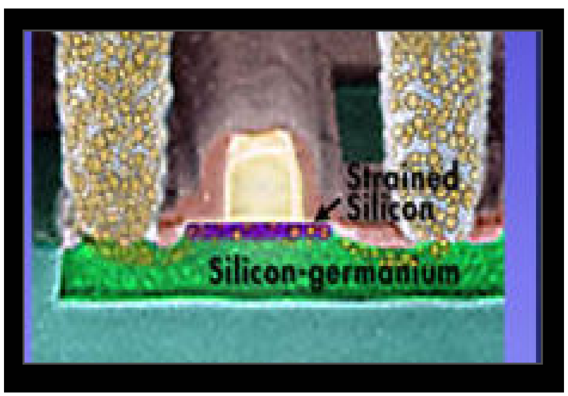

2001 Dual Stress Liner

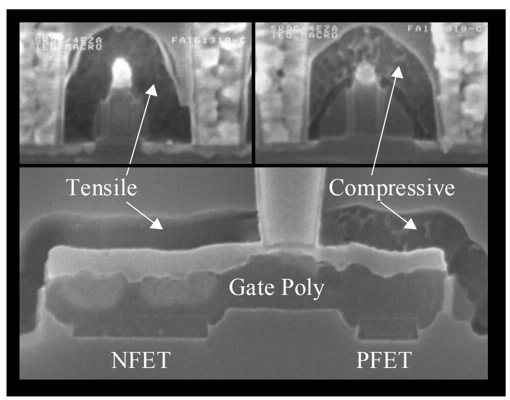

2003 Self-Aligned
Carbon Nanotube FET

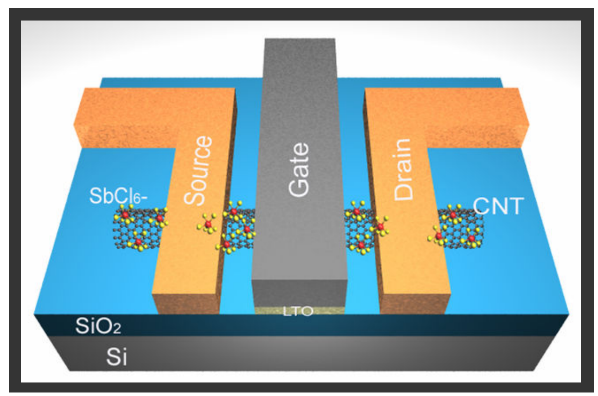

ONLY Innovation Can Make It Work

	2004	2007	2010	2013	2016	2020
Physical Gate	37 nm	25 nm	18 nm	13 nm	9 nm	6 nm
Keff	3.0	2.4	1.9	1.5	1.3	<1.3

Device Roadmap

Ultrathin SOI

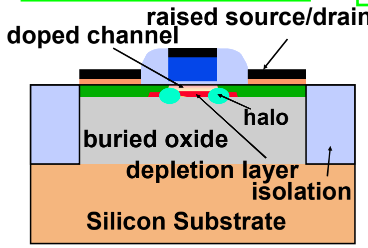

High k gate dielectric

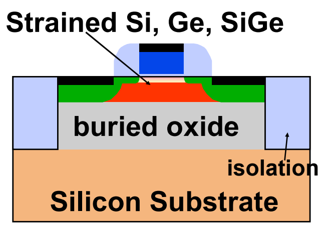

Double-Gate CMOS

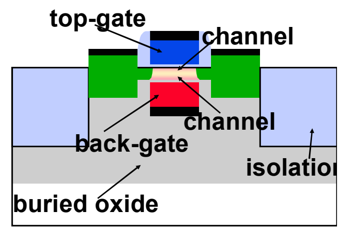

FinFET

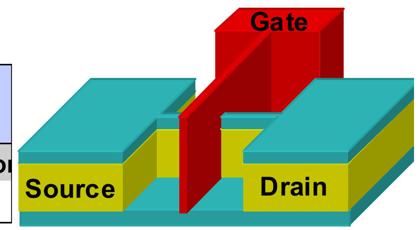

Interconnect Roadmap

Low k

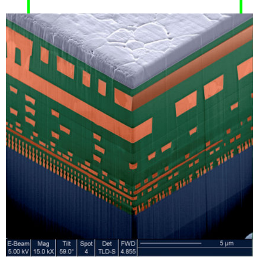

Ultra Low k

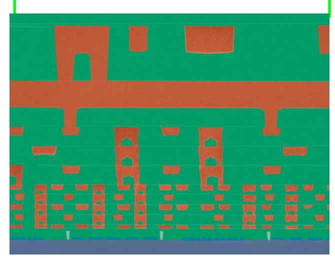

Porogens

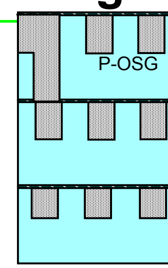

AirGap

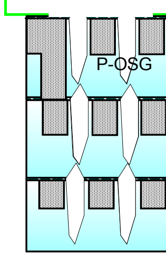

Holistic Design:
A “New” Paradigm in Value Creation

Innovation from Atoms to Software

- **The simultaneous optimization of:**

Materials, Devices, Circuits, Cores, Chips, System Architecture, System Assets and System Software

- **Provides the most effective means to optimize the value of IT offerings to the end user**

Note:

Execution relies upon the seamless integration of IBM skills from across the spectrum

Holistic Design Delivered; Power 5+ and Power 6

"Despite all the challenges cited here, we are not remotely near the end of this journey, e.g. Don't sell technologists short"

Shipped — Hardware

2001 — POWER4™

Industry 1st dual core

180 nm

1.0-1.3 GHz Core | 1.0-1.3 GHz Core
Shared L2
Distributed Switch

- Chip multiprocessing
  - Distributed switch
  - Shared L2
- Dynamic LPARs (16)

2002-4 — POWER4+™

130 nm

1.2-1.9 GHz Core | 1.2-1.9 GHz Core
Shared L2
Distributed Switch

- Reduced size
- Lower power
- Larger L2
- More LPARs (32)

2004 — POWER5™

130 nm

1.5-1.9 GHz Core | 1.5-1.9 GHz Core
Shared L2
Distributed Switch

- Micro-Partitioning
- Virtual storage, virtual ethernet
- Simultaneous multi-threading
- Dynamic firmware updates
- Enhanced scalability
- High throughput performance
- Enhanced storage subsystem
- **Quad Core Module**

POWER5+™

90 nm

> GHz Core | > GHz Core
Shared L2
Distributed Switch

>2Ghz

POWER6™

65 nm

Ultra-high frequency cores
L2 caches
Advanced System Features

4-5Ghz

Optimizations announced
- 4-5Ghz. @ reduced power density
- Buffered H clock tree
- Transmission line distribution
- High speed data cache
- Low power/low latency static circuits
- 13FO4 design
- Power 6 systems in mid-'07.

Microprocessor Forum 2004

* All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Holistic Design: Virtualization as a System Strategy

Virtual Resources

- Proxies for physical resources: same interfaces/functions, different attributes.
- May be part of a physical resource or multiple physical resources.

Virtualization

- Creates virtual resources and "maps" them to physical resources.
- Primarily accomplished with software and/or firmware.

Physical Resources

- Hardware components with architected interfaces/functions.
- May be centralized or distributed.
- Examples: cpu, memory, storage, networks, servers, ...

- **Separates presentation of resources to users from actual physical resources**
- **Aggregates pools of resources for allocation to users as virtual resources**

Innovative Design Principles

Integrated Processors

- Cell

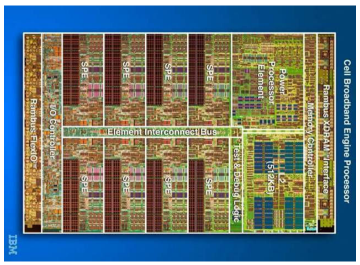

- Secure Blue

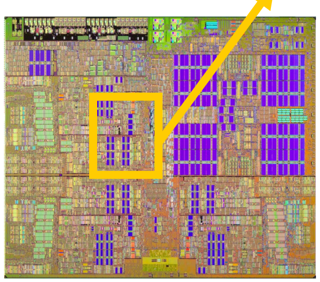

- Blue Gene

"Scale-out" at the Chip Level

- IBM's Partnership with Sony & Toshiba has yielded a revolutionary chip architecture known as Cell

SONY

IBM

TOSHIBA

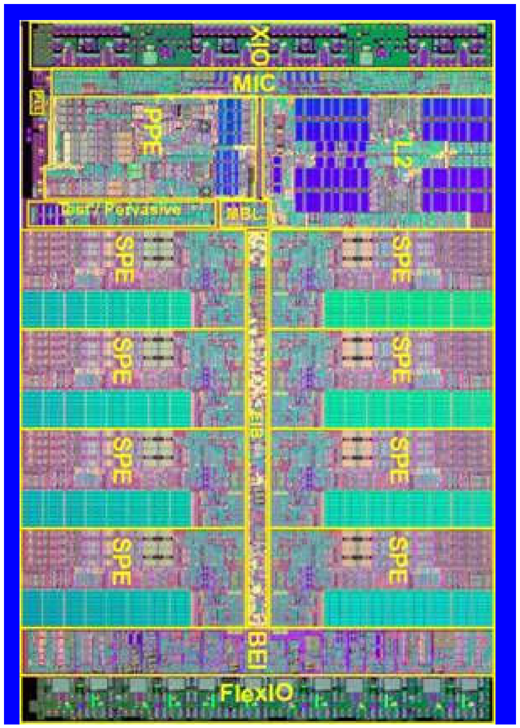

- 234M Transistors
- 64 Bit Power Processor
- 8 Synergistic Processor Elements
- 256 Billion Floating Point Operations per second (S.P.)

Investment in Capital + aggregation of Intellectual Capital that will drive future progress

BlueGene/L: "Scale-Out" at the System Level

- **BlueGene/L**: **1/100th the physical size** (320 vs 32,500 square feet) consumes **1/28th the power** (216KW vs 6,000KW) - compared to Earth Simulator

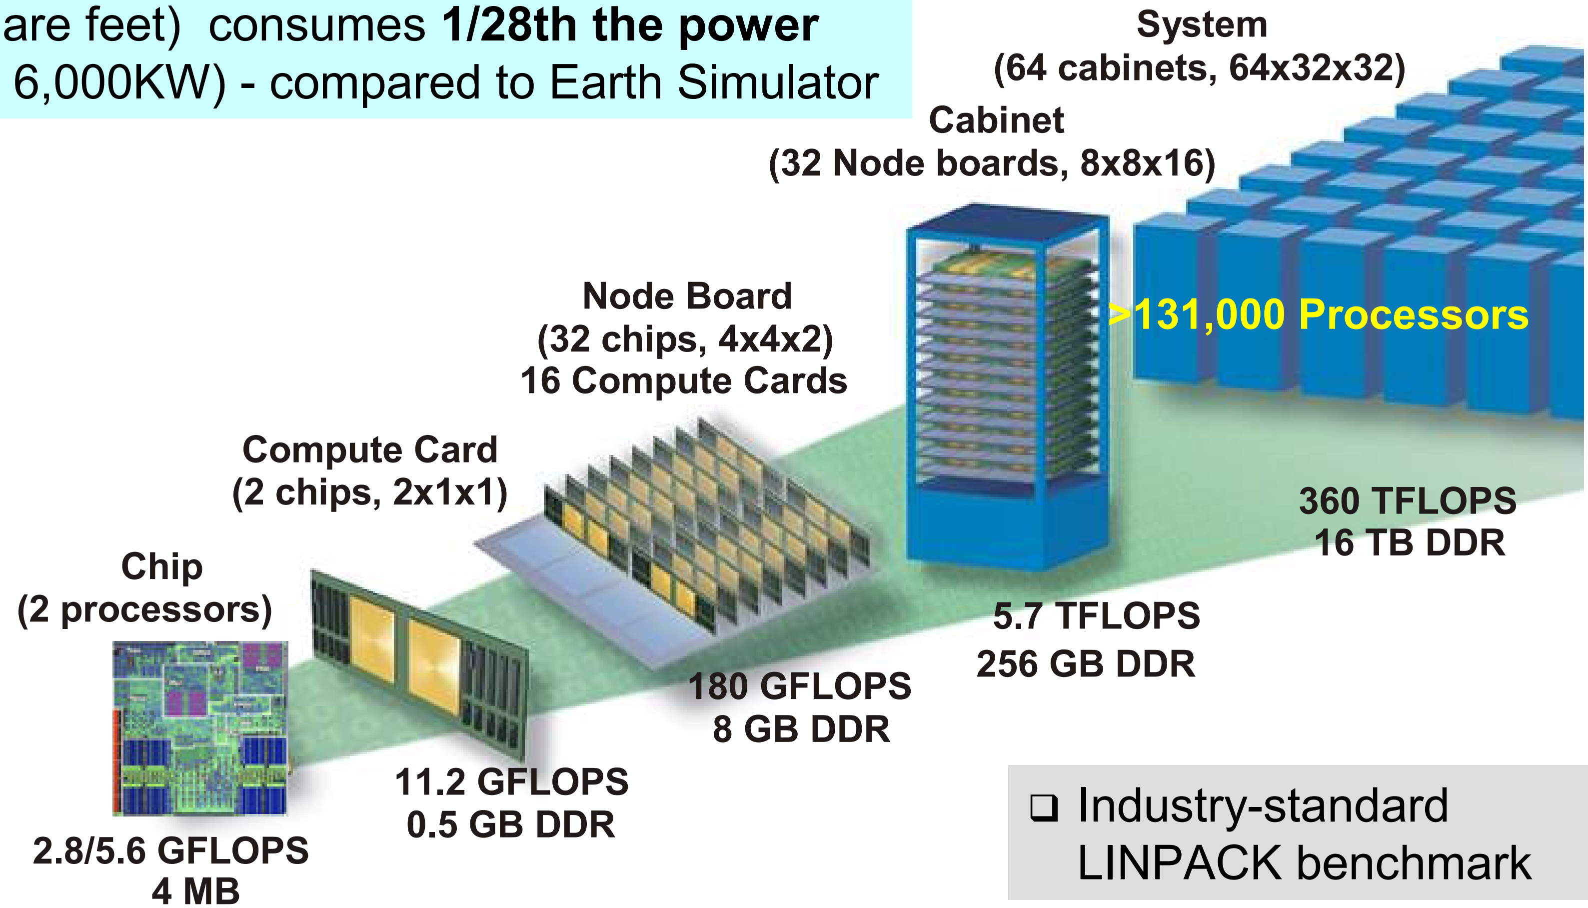

- Industry-standard LINPACK benchmark

- Attained a sustained performance of **<u>70.72 Teraflops</u>**
  - eclipsing 3 year old top mark of **35.86 Teraflops** - Japanese Earth Simulator
  - recent mark of **42.7 Teraflops** at the NASA's Ames research center

Innovative Design Principles

Extended systems

Mainframe

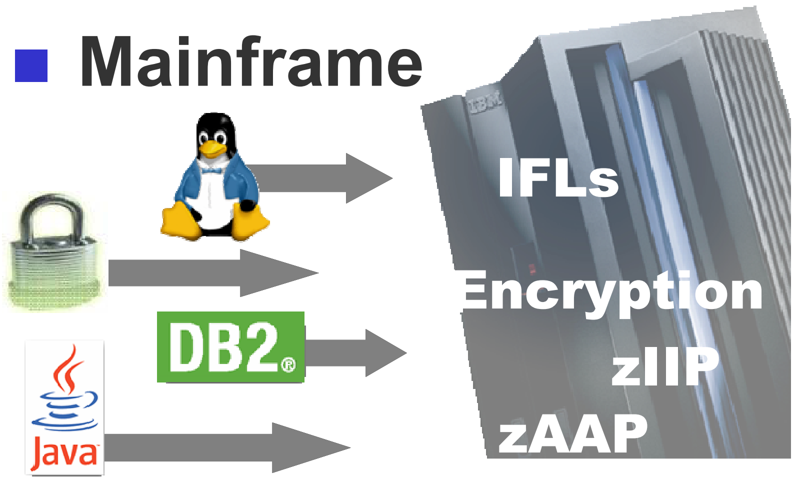

Power Systems

Blade Centers

x86
1. Journaling
2. Store Apps
3. Symbol WLAN
4. Cache
5. PICS / StoreTV
6. Int Voice Resp
7. Voice over IP
8. Point of Sale

Power
9. Pharmacy
10. Merchandizing
11. NIM / Deploy

12. iSCSI
13. iSCSI
14. iSCSI

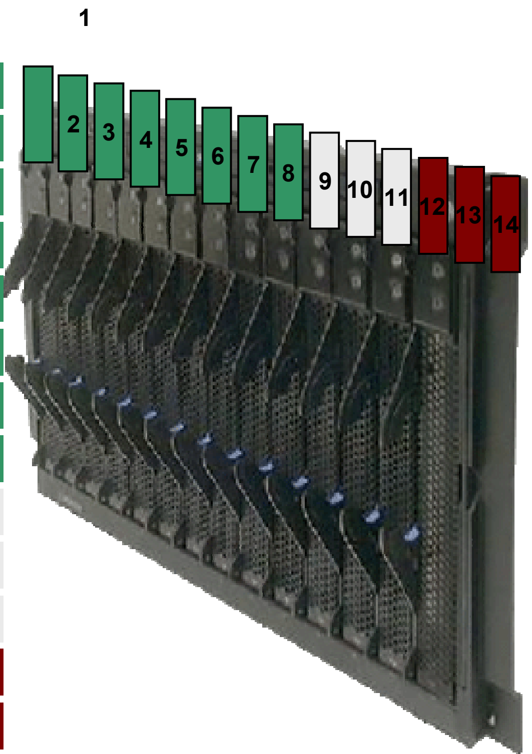

BladeCenter Store in a Box Solution

Innovative Design Principles

- Integrated Processors
- Extended Systems
- **Architected for Collaboration**

Power.org™

J2EE

the globus consortium

Server Blade Summit

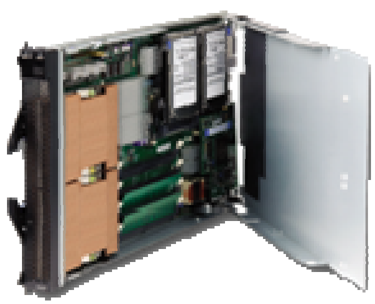

Blade.org

The Apache

IBM Virtualization Engine

Xen

The New Economics of IT; A Paradigm Shift

- **Power and cooling spend will exceed new server spending**

2000 – Raw processing "horsepower" is the primary goal, while the infrastructure to support it is assumed ready

2006 – Raw processing "horsepower" is a given, but the infrastructure to support deployment is a limiting factor

Three Cooling Challenges

1. The System
2. The Rack
3. The Data Center

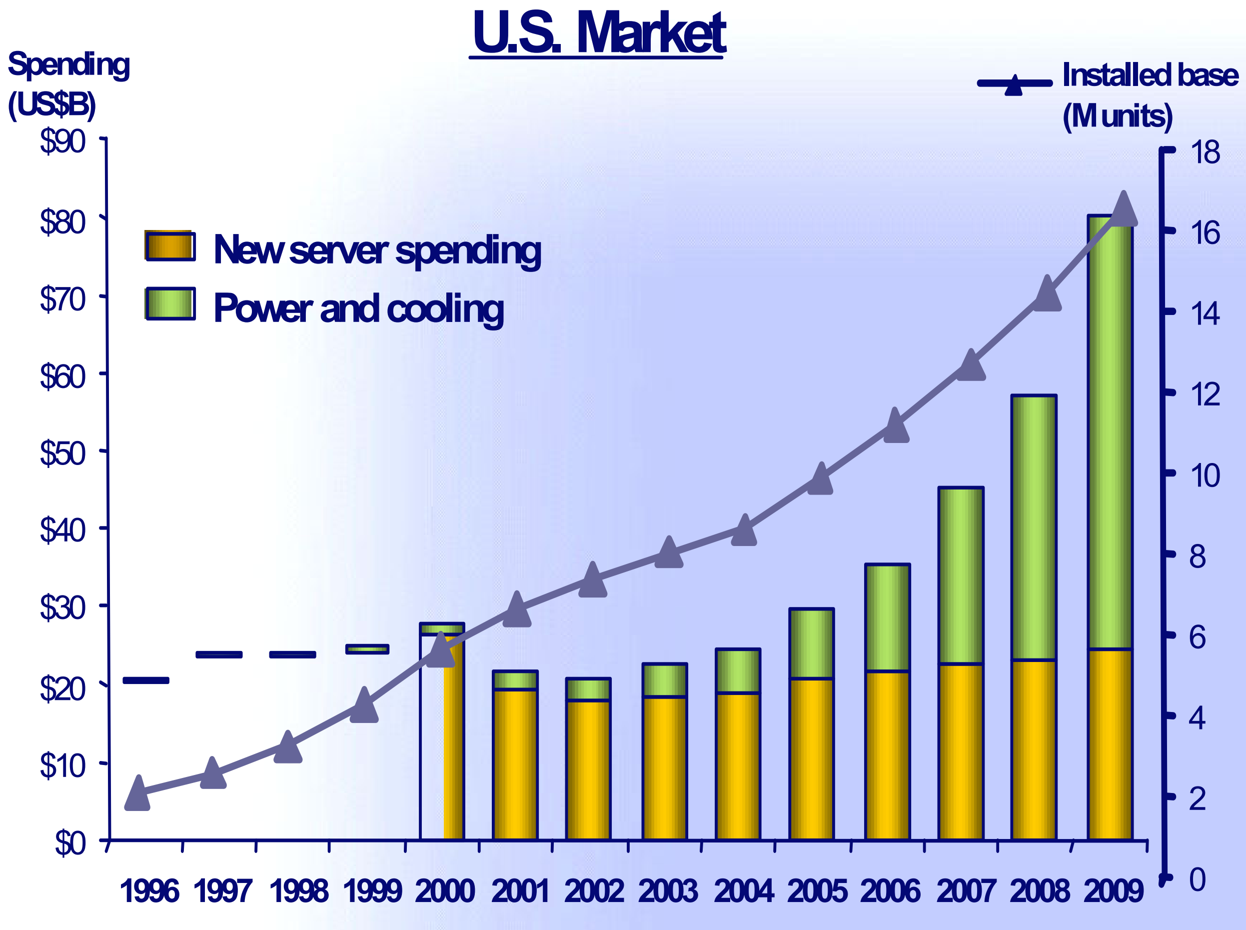

Source IDC: 2006, Document # 201722, "The Impact Of Power and Cooling On Data Center Infrastructure", John Humphreys, Jed Scaramella

The System-Issue #1
Thermal and Power Advances

- **Advanced designs allow IBM to pack more function into a server than our competitors**
  - More memory, PCI bus, and better redundancy in our 1U and 2U servers
  - More servers in a chassis with BladeCenter
- **IBM builds better efficiency into its systems**
  - Shared system infrastructure components (90%)
  - Utilization of “power factor corrected” power supplies - deliver high energy efficiency to reduce wasteful electrical usage & heat output
  - Use of lower power components to stretch the available power budget of the system
  - Calibrated Vector Cooling technology - dual paths of air to each component: improved uptime, longevity, & air movement

Efficient power supply conversion AC to DC

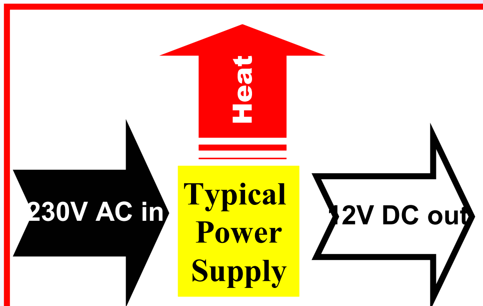

Example: 1000W AC in at 70% efficiency = **700W DC** output

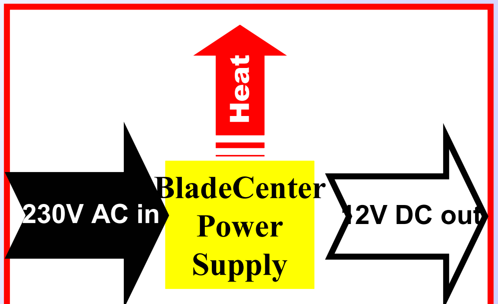

Example: 1000W AC in at 91% efficiency = **910W DC** output

Understanding Where Costs Hide is Critical

- **IBM's CoolBlue Initiative includes Rear Door Heat eXchanger which can remove over 50% of a rack's heat output**
  – No new fans or electricity needed.
  – Attaches to back of rack (adds 5") No rearrangement of datacenter
  – Cost effective; 1KW cooling = $286

➲ **Rear Door Heat eXchanger adds cooling capacity at ~1/4 of the cost of traditional methods**

Normal Operation of Server

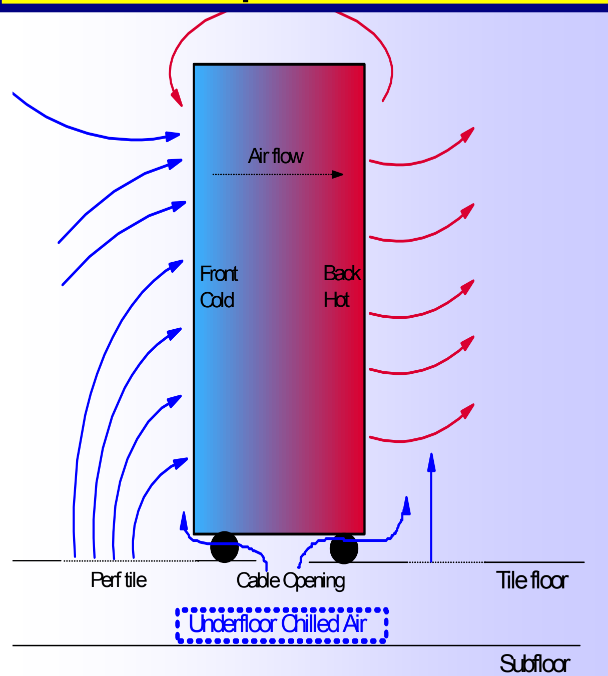

Improved Operation of Server

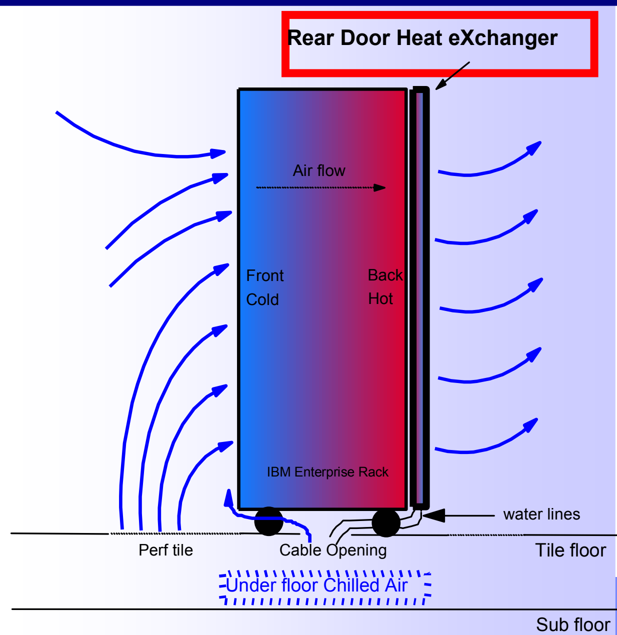

Rear Door Heat eXchanger

Typical Data Center Power Distribution and Energy Savings w/Water

Air Cooled Server Racks

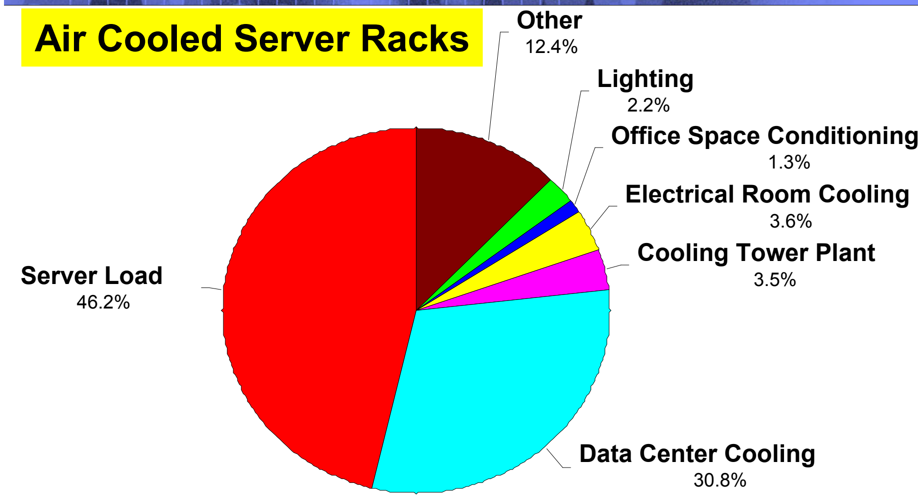

Description	Electricity Consumption(kW) with Air Cooled Servers	Electricity Consumption(kW) with Water Cooled Servers
Other	402	402
Lighting	73	73
Office Space Conditi	41	41
Electrical Room Coo	118	118
Cooling Tower Plant	114	114
Data Center Cooling	1000	418
Server Load	1500	1500
Total Load	3248	2666

Water Cooled Server Racks

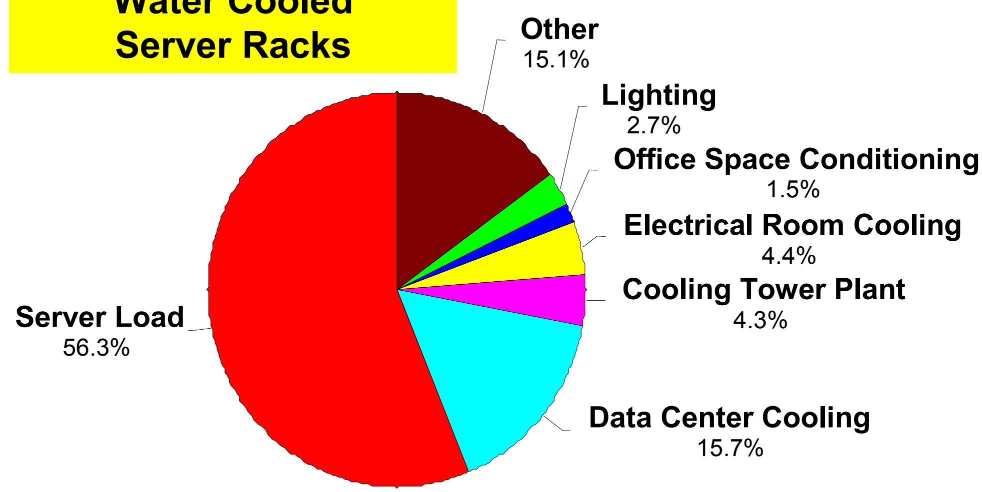

Savings of ½ million dollars/year at $0.10 / kw-hr

IBM "Smart Tools" – Thermal Diagnostics

- **IBM Power Configurator and Power Executive**
  - **Tools to deliver better sizing information about our clients solutions**
  - **A powerful software suite designed to give users precise information as to power consumption, enabling better decision making**

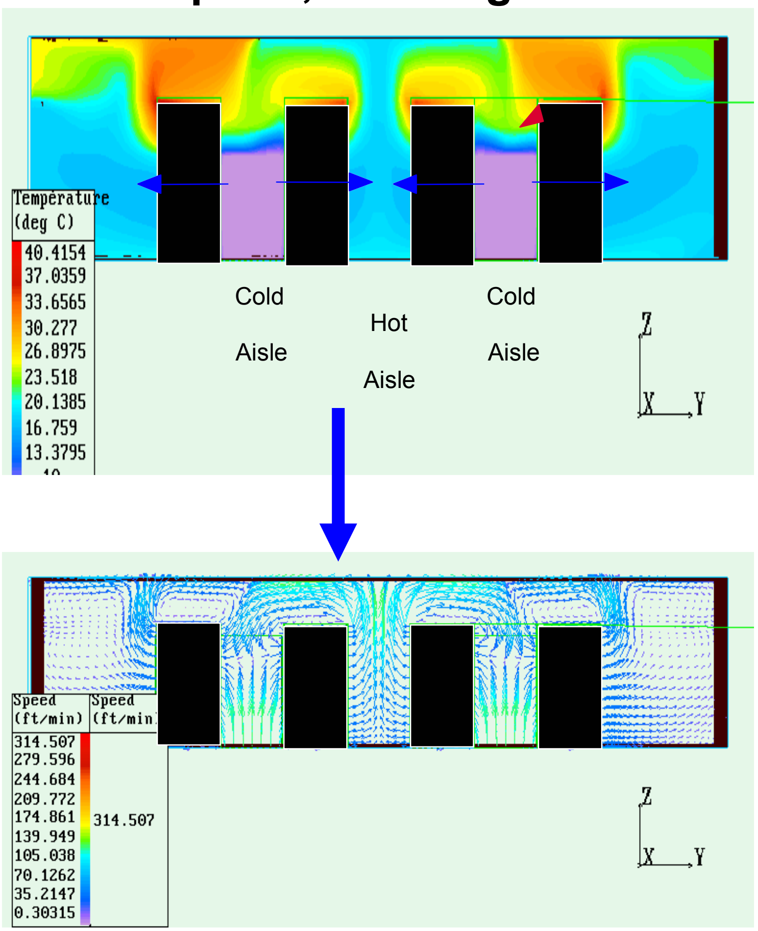

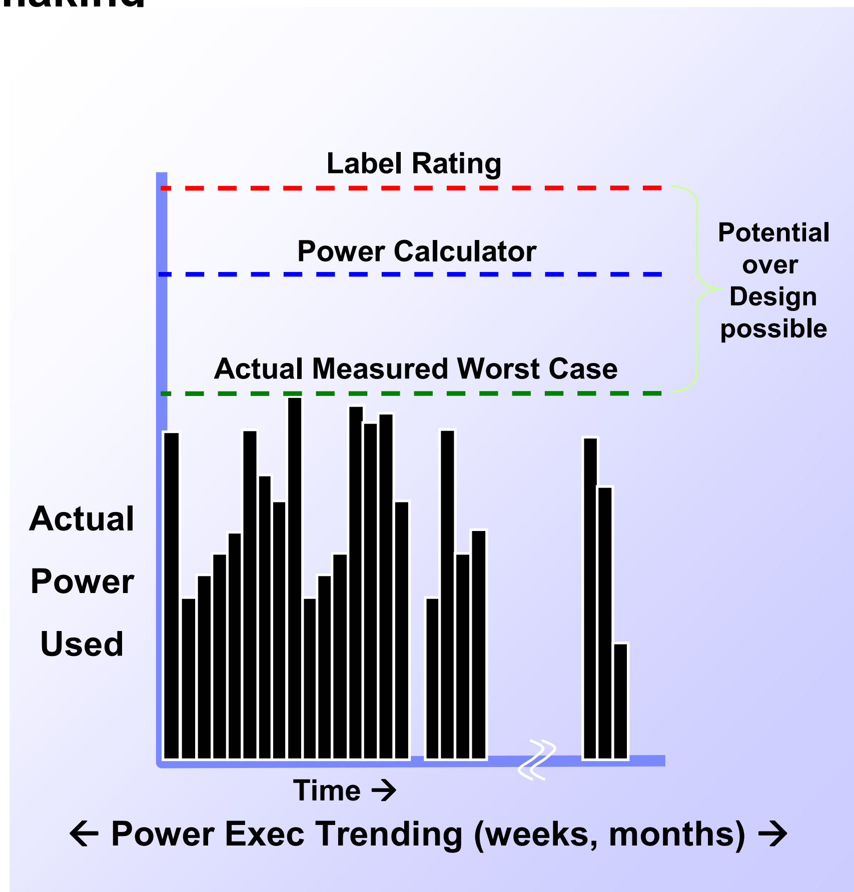

A Key Factor Impacting Innovation: *COST!*

Innovations are EXPENSIVE!

- **Lithography**
  - Dry to Immersion to EUV(?)
- **Transistors**
  - From Si → strained Si, high-K, double Gate, FinFet.
- **Interconnects**
  - From Al and Cu → LowK, UltraLowK, Porigins, AirGap
- **Manufacturing**
  - From traditional → on demand

Semiconductor Industry's Challenge

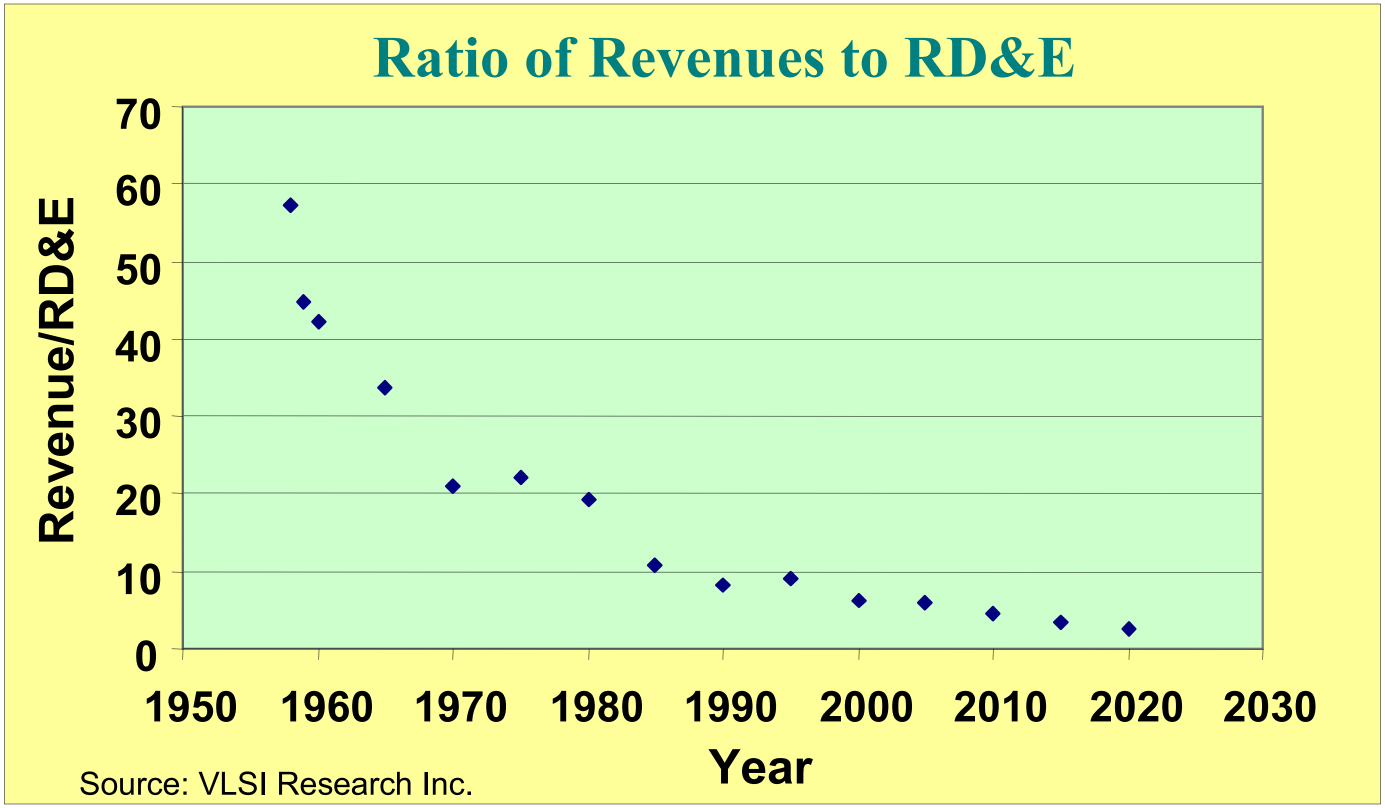

- **This is the industry trend as of today.**
- **This is NOT sustainable, and will/must NOT happen.**

Globalization of Semiconductor R&D is Already a Pervasive Strategy

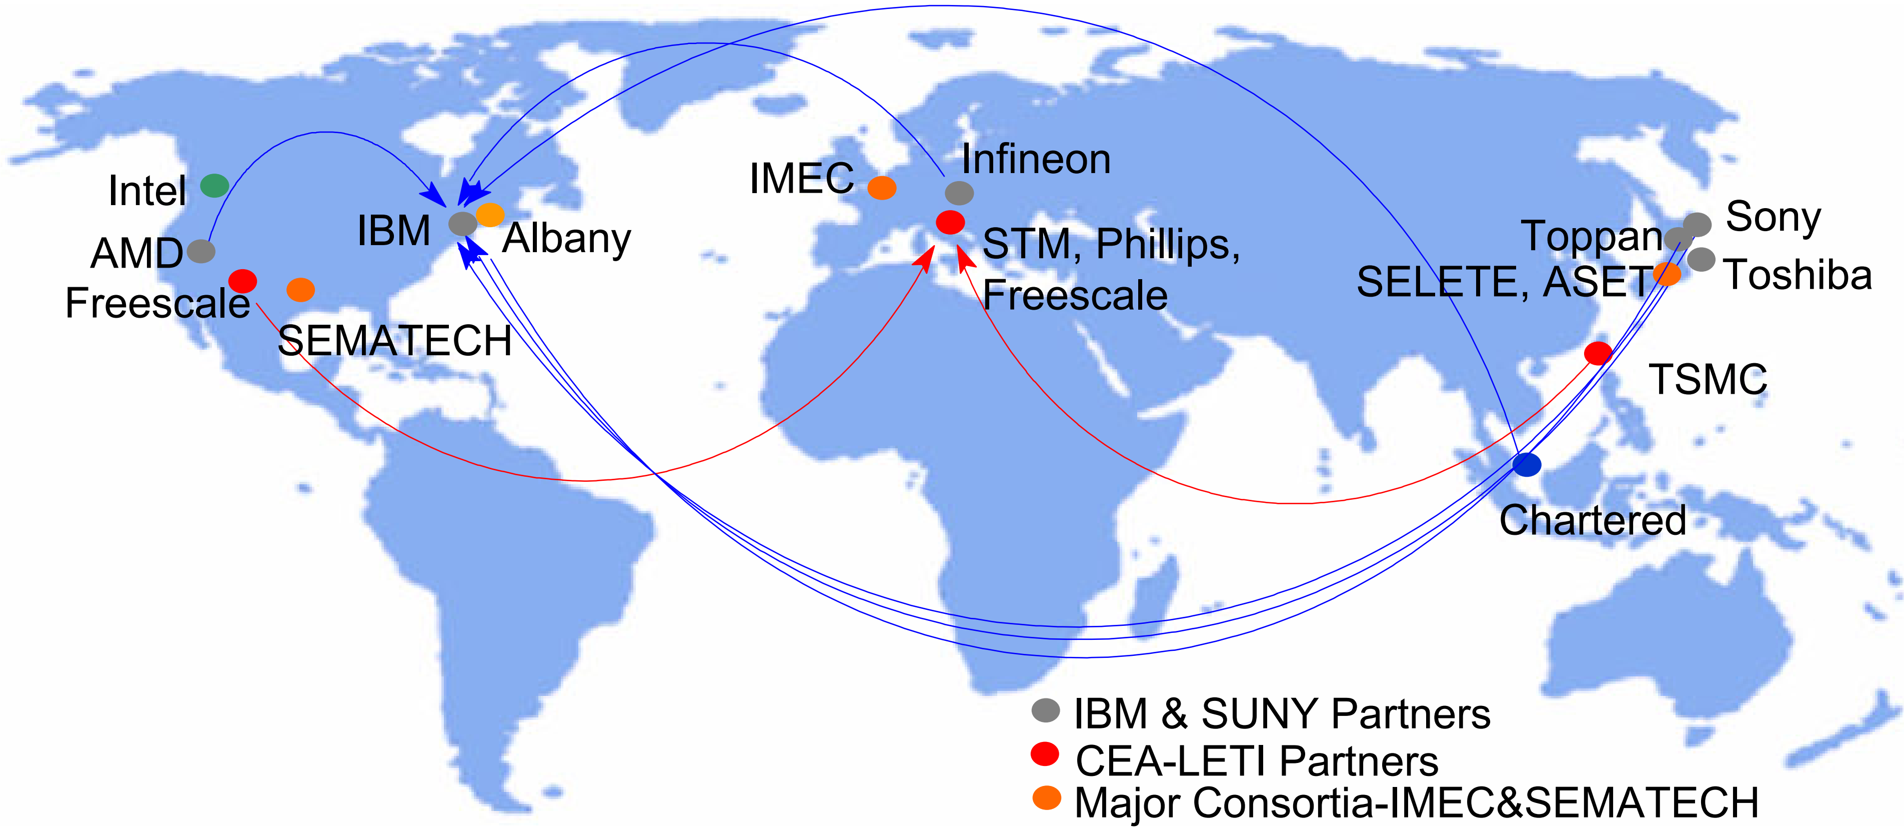

There is a global recognition of challenges facing semiconductor industry, having already driven a its dramatic consolidation

=> Pre-competitive cooperation

The New Levers

- **Innovation & integration delivers what scaling can't**
  - High frequency systems at high functionality have been demonstrated via Holistic Design
  - Thermal issues will become ever more severe over time, and fundamental innovations are required to provide a route forward

- **System solutions optimized via Holistic design will ultimately dominate progress in information technology**
  - One lever: use of many levers from across IBM & Partners in balanced approach

- **Fiscal reality is driving industry towards consolidation around Innovation Networks, the new financial "lever"**
  - This requires globalization of technology ecosystems

- **Collaborative innovation**
  a key strategic imperative
  - 1st emergence of open standards & technology ecosystems, enabling collaborative innovation in what was a highly proprietary endeavor